FACT SHEET

HP2C – an initiative for more computing efficiency

Model simulations have become a key element of modern scientific methodology. As the models become increasingly complex, however, the algorithms need to be improved and the supercomputers rendered more powerful. Switzerland is tackling the software problem with the HP2C project.

Top-notch research needs supercomputers, like those available to Swiss researchers at CSCS (Swiss National Supercomputing Centre) in Lugano. With the increasing complexity of computational models, high-performance computing has become a key tool of modern research. Computer models for solving scientific problems are therefore an established component of scientific method alongside experiments and theory. The simulations, which are created with supercomputers that work massively parallel nowadays, provide completely new insights in science and boost the creativity of engineers. They visualise what remains invisible to the naked eye, even with high-resolution microscopes or telescopes. Simulations help us to discover unknown materials with completely new properties and functionalities, for instance.

Unbridled increase in performance

Twenty years ago, supercomputers achieved the same performance level as today's laptops: back then, a high-performance computer in the so-called gigaflops performance class carried out a few billion calculations per second; today's supercomputers in the petaflops class perform several quadrillion – that's several million times a billion. As conventional chip technologies are stretched to their limits, computers from the petaflops performance class have to be constructed from several thousand to tens of thousands of well-networked processors, which in turn results in a massive increase in energy consumption. The most powerful supercomputers in the world consume about ten megawatts of electricity. One megawatt costs a computing centre around CHF 1.5 million a year.

Experts therefore agree that supercomputers need to be utilised more effectively. In concrete terms, this means that application programmes for massively parallel computers have to be rewritten. Consequently, Thomas Schulthess, Director of CSCS, teamed up with Piero Martinoli, President of Università della Svizzera italiana, against the backdrop of the national High-Performance Computing and Networking Strategy (HPCN Strategy) in 2009 to launch the High-Performance and High-Productivity Computing Initiative (HP2C). Under HP2C, developers of application software for scientific simulations from different disciplines collaborate with mathematicians, applied mathematicians and computer scientists in an interdisciplinary approach.

Improving algorithms and software

A total of thirteen project groups from chemistry, materials science, physics, astrophysics, medicine, anthropology, and the climate and geosciences are involved in the HP2C-project. They all work with highly complex, multifaceted models and, like many computer-based sciences, face the problem that the codes they have been using on the hardware of the massively parallel computers thus far and their chips with an increasing number of processor cores no longer perform efficiently.

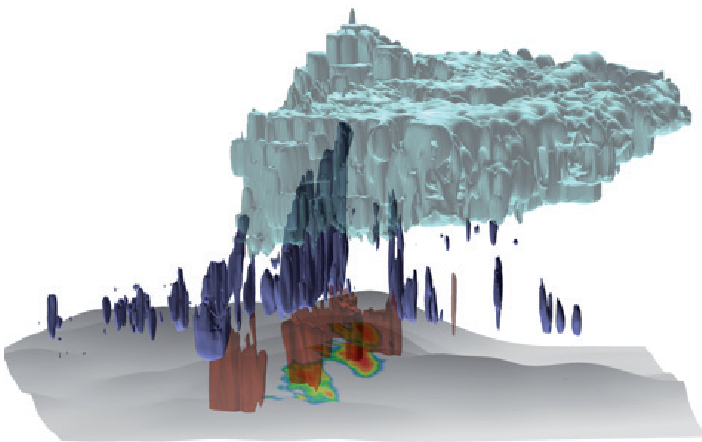

The cloud-resolving models used by the ETH Zurich climate research team are based on grid point distances of around 1 km. They can simulate summer storm clouds, as shown here, in detail. Clouds of water droplets are shown in dark blue, whereas falling droplets are marked in red. Rainbow colours indicate the resulting amount of precipitation (red=high, blue=low). *(Image: Wolfgang Langhans, ETH Zurich)*

The aim is for the application software in computer simulations to capitalise on the potential computational power of modern computer architectures. Consequently, the computing methods are being adapted under HP2C in such a way as to make optimum use of the processor hardware. Moreover, the simulation programmes are being structured in such a way as to minimise the movements of the data on the processor, between the processor and main memory, and between the processors of the massively parallel supercomputer.

Testing the new hardware and software

The codes for the different questions have to be modified individually. In order to test them on state-of-the-art hardware, the supercomputer "Tödi", a Cray XK6, was installed at the CSCS in October 2011. "Tödi" is a massively parallel supercomputer based on a hybrid technology. The technology uses a mixture of conventional processors with graphics processors – according to experts, a promising approach for future technology in high-performance computing. Unlike conventional computer processors, graphics processors run hundreds or even thousands of calculations in parallel, which not only helps researchers to solve their problems more quickly, but also increases energy efficiency considerably. Graphics processors are up to ten times more powerful than conventional processors, but only consume twice the electricity.

Laurent Villard, chief initiator of the HP2C project "Gyrokinetics":

Laurent Villard, Professor of plasma physics at EPF Lausanne.

Why did you apply for an HP2C-Project?

Turbulence simulations in magnetically confined plasmas are highly challenging as they involve disparate time and length scales. It is therefore essential that our codes are upgraded, and if necessary re-factored, in order to tap the capabilities of the most advanced HPC platforms. The development of our codes must be placed in a long-term perspective: more and more physical effects are being included, requiring ever increasing computing resources. Thus, it is crucial to keep up with the most recent developments and have a prospective view on HPC issues.

What's the main benefit for you and your research team of HP2C?

The nice thing about HP2C is that it puts financial resources in the research institutes that conduct the specific scientific application, and at the same time develop links with a team of high-level specialists at CSCS. Thanks to HP2C we could further improve the parallel scalability of our codes. In the most demanding cases (several billion grid points) the performance could be increased by a factor of two. Moreover, it has stimulated us to investigate new avenues for future code developments. One of the "fringe" (but non-negligible!) benefits of HP2C is that we could secure access to a petaflops-class HPC platform since early 2012.

How would you complete this sentence: HP2C is for Switzerland…

… a unique, well-thought, well-managed opportunity to keep us at the forefront of science in several areas that require ever increasing computing power.

Project examples

The following examples demonstrate the direct impact the success of HP2C projects has on our everyday lives:

Furthering earthquake research

The idea behind the project “Petaquake” is to clear up key questions in the geosciences, such as how exactly the Earth’s interior is structured, which processes take place there and how earthquakes develop. A glimpse inside the Earth (akin to computer imaging in the field of medicine) could provide the answers. The codes and algorithms developed in the project should enable high-resolution tomographic images of processes in the Earth’s interior to be calculated on a ten-kilometre scale. These images could help reduce existing uncertainties in risk assessment and improve hazard maps. Petquake has a direct practical benefit for earthquake-prone Switzerland and Europe.

Petaquake – Large-Scale Parallel Nonlinear Optimization for High-Resolution 3D-Seismic Imaging; Olaf Schenk, Professor at the Università della Svizzera italiana

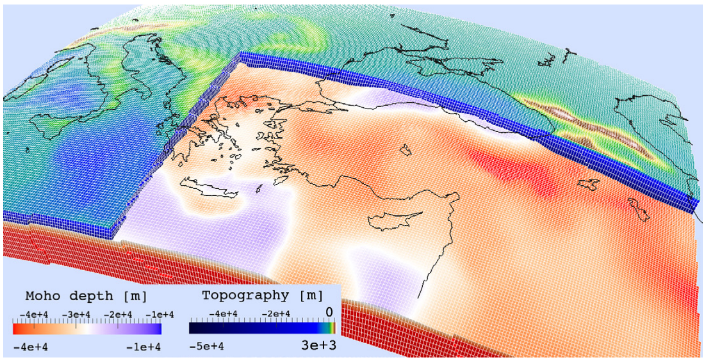

A realistic model of seismic wave propagation in Europe. The mesh comprises 2 billion grid points, surface topography, and three-dimensional structure of the subsurface. It was developed by seismologists of ETH Zurich on supercomputers of CSCS for simulating earthquake scenarios. Such realistic subsurface structures are essential for quantifying earthquake risks, for example for the area of Istanbul in Turkey. The models can then be used for assessing critical local building structures. *(Image: Rietmann et al., Supercomputing 2012, accepted)*

Solving energy problems with the supercomputer

Simulations can aid research on drugs or alternative energy sources – novel fuels or solar cells. Special software applications enable new materials and their highly complex molecular properties to be simulated. One of the most important codes for such simulations is CP2K (the Car-Parrinello 2000 project), which can be used to simulate large, complex chemical systems, such as interfaces with material transitions from solid to liquid or liquid to gas. Solar cells such as Grätzel cells are a prime example of this kind of interface: a solid body on one side and a solvent on the other, with a dye in between. This electrochemical dye solar cell works according to a similar principle to photosynthesis and currently achieves an efficiency level of 12.3 percent. In order to be able to build the cell even more efficiently, the researchers need to know how it works as precisely as possible. This is where simulations come in. The better the solar cells are understood based on the simulations, the more efficiently one can experiment with them in the lab and improve their level of efficiency.

CP2K uses many different algorithms and has complex numerical structures that, among other things, need to be represented in the new computer architectures within the scope of HP2C.

CP2K – New Frontiers in Ab Initio Molecular Dynamics; Jürg Hutter, Professor at the University of Zurich

Agent-based simulations – learning to understand behaviour

Between 20 000 and 200 000 years before our time, at least two species of hominid cohabited with Neanderthal and modern man on Earth. While modern and Neanderthal man existed separately from each other in Africa and Europe respectively, Homo sapiens and Homo neanderthalensis lived side by side in the Middle East for around 100 000 years. Why the Neanderthals suddenly became extinct over 20 000 years ago, however, remains a mystery to this day.

Using a so-called agent-based model, researchers are looking to simulate different behaviour scenarios between Neanderthals and man, taking the fauna, climate conditions, topography and vegetation into account, in order to explore the possible causes of Neanderthal extinction more precisely – a complex modelling technique in which every individual, equipped with specific characteristics, is represented as a discrete object. This kind of model is becoming increasingly important in simulating the behaviour of people, such as in the event of a mass panic, for instance.

Neanderthal Extinction – Individual-Based Modelling of Humans under Climate Stress; Christoph P. E. Zollikofer, Professor at the University of Zurich

Electricity through nuclear fusion

With the gyrokinetics project in HP2C, plasma physicists from EPF Lausanne are pursuing the goal of developing new numerical codes for the simulation of turbulent transport in plasma with magnetic confinement. In doing so, they are making a major contribution to ITER, the largest international project in the field of thermonuclear fusion research. ITER (Latin for "the way") is expected to provide evidence that generating energy from the nuclear fusion of deuterium and tritium is physically and technologically feasible. ITER is based on the tokamak principle and will confine the hot fusion plasma using magnetic fields generated by superconductive coils. Controlling plasma turbulence is important as it is responsible for increased heat transport and thus limits the reactor's energy confinement period. Laurent Villard and his team's work in HP2C has attracted much attention all over the world, which is reflected in the fact that a quarter of the entire computation time of a 1.5-petaflops supercomputer earmarked especially for fusion research in Japan was approved to test their newly developed codes.

GYROKINETICS – Advanced Gyrokinetic Numerical Simulations of Turbulence in Fusion Plasmas; Laurent Villard, Professor at EPF Lausanne

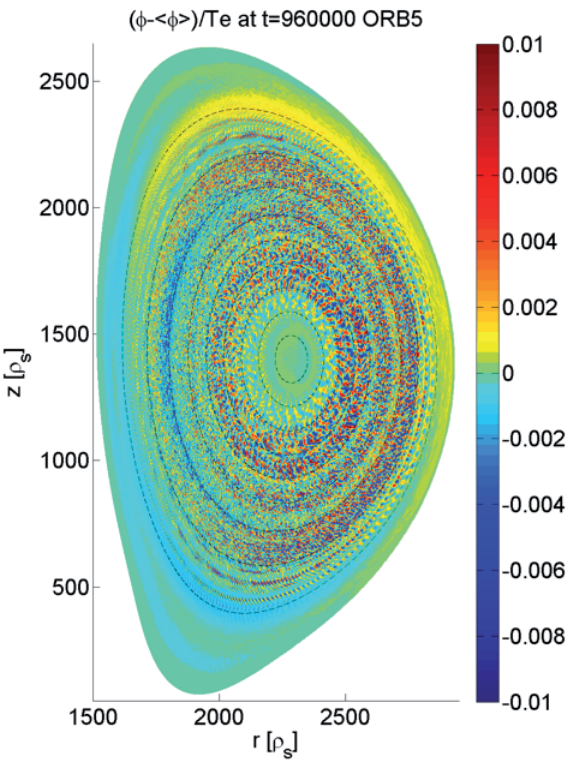

The figure shows contours of the density fluctuations in an ITER turbulent plasma. Simulation performed with the ORB5 code on the IFERC-CSC HELIOS platform (1.5 petaflops), using 1 billion grid points and 2 billion numerical particles. "The HP2C project has been very successful for us", says Laurent Villard. "The ORB5 code performance and parallel scalability was substantially improved and enabled us to perform large scale turbulence simulations that were practically unfeasible before. It gave us access to petascale computing." *(Image: Gyrokinetic group Laurent Villard, EPF Lausanne).*

HP2C High Performance and High Productivity Computing

In 2008 Thomas Schulthess formulated a project for the Rectors' Conference of Swiss Universities (CRUS). Its aim was to reinforce the collaboration between the CSCS and USI, and tackle basic problems in supercomputing in a structured fashion. Originally, the HP2C project was intended to be smaller and entail an investment of approximately CHF four million. However, the Swiss University Conference suggested expanding the platform throughout Switzerland and contributes approximately eighty per cent of the funding – around CHF 14.5 million. HP2C is a central element of the national strategy for high-performance computing funded by the federal government since 2009.

www.hp2c.ch

Via Trevano 131
6900 Lugano
Switzerland

Tel +41 (0)91 610 82 11
Fax +41 (0)91 610 82 82
www.cscs.ch